(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 087 635 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
28.03.2001 Bulletin 2001/13

(51) Int Cl.7: **H04Q 11/00**, H04Q 11/04

(21) Application number: **00308473.8**

(22) Date of filing: **27.09.2000**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **27.09.1999 CA 283627**

(71) Applicant: **NORTEL NETWORKS CORPORATION**
Montreal Quebec H2Y 3Y4 (CA)

(72) Inventor: **Beshai, Maged E.**
Stittsville, Ontario K2S 1E2 (CA)

(74) Representative: **Ertl, Nicholas Justin**
Elkington and Fife,
Prospect House,
8 Pembroke Road
Sevenoaks, Kent TN13 1XR (GB)

(54) High-capacity WDM-TDM packet switch

(57) A self-configuring distributed packet switch(20) that operates in wavelength division multiplexed (WDM) and time division multiplexed (TDM) modes is described. The switch comprises a distributed channel switching core (26), that includes core modules (34) which are respectively connected by a plurality of channels to a plurality of high-capacity packet switch edge modules (22,24). Each core module operates independently to schedule paths between edge modules, and reconfigures the paths in response to dynamic changes in data traffic loads reported by the edge modules. Reconfiguration timing between the packet switch modules and the channel switch core modules is performed to keep reconfiguration guard time minimized. The advantage is a high-capacity, load-adaptive, self-configuring switch that can be distributed to serve a large geographical area and can be scaled to hundreds of Tera bits per second to support applications that require very high bandwidth and a guaranteed quality of service.

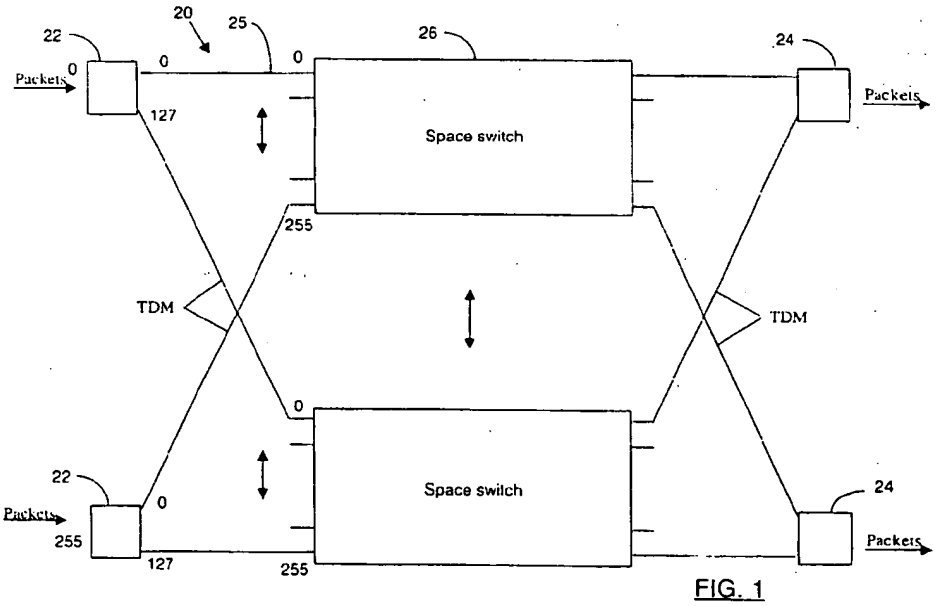

FIG. 1

EP 1 087 635 A2

Printed by Jouve. 75001 PARIS (FR)

Description

TECHNICAL FIELD

[0001] This invention relates generally to the field of data packet switching and, in particular, to a distributed very high-capacity switch having edge modules that operate in packet switching mode and core modules that operate in circuit switching mode, the core modules switching payload traffic between the edge modules using wavelength division multiplexing (WDM) and time division multiplexing (TDM).

BACKGROUND OF THE INVENTION

[0002] Introduction of the Internet to the general public and the exponential increase in its use has focused attention on high speed backbone networks and switches capable of delivering large volumes of data at very high rates. In addition to the demand for higher transfer rates, many service applications are being developed, or are contemplated, which require guaranteed grade of service and data delivery at guaranteed quality of service. To date, efforts to grow the capacity of the Internet have largely been focused on expanding the capacity and improving the performance of legacy network structures and protocols. Many of the legacy network structures are, however, difficult to scale into very high-capacity networks. In addition, many legacy network protocols do not provide grade of service or quality of service guarantees.

[0003] Nonetheless, high capacity switches are known in the prior art. Prior art high capacity switches are commonly constructed as a multi-stage, usually three-stage, architecture in which ingress modules communicate with egress modules through a switch core stage. The transfer of data from the ingress modules to the egress modules must be carefully coordinated to prevent contention and to maximize the throughput of the switch. Within the switch, the control may be distributed or centralized. A centralized controller must receive traffic state information from each of the ingress modules. Each ingress module reports the volume of waiting traffic destined to each of the egress modules. The centralized controller therefore receives traffic information related to traffic volume from each of the ingress modules. If, in addition, the controller is made aware of the class of service distinctions among the waiting traffic, the amount of traffic information increases proportionally. Increasing the amount of traffic information increases the number of control variables and results in increasing the computational effort required to allocate the ingress/egress capacity and to schedule its usage. Consequently, it is desirable to keep the centralized controller unaware of the class of service distinctions while providing a means of taking the class of service distinctions into account during the ingress/egress transfer control process.

[0004] This is accomplished in a rate-controlled multi-class high-capacity packet switch described in Applicant's copending European Patent Application No. 00300700.2 which was filed on January 31, 2000. Although the switch described in that patent application is adapted to switch variable sized packets at very high speeds while providing grade-of-service and quality-of-service control, there still exists a need for a distributed switch that can form the core of a powerful high-capacity, high-performance network that is adapted to provide wide geographical coverage with end-to-end capacity that scales to hundreds of Tera bits per second (Tbs), while providing grade of service and quality of service controls.

[0005] A further challenge in providing a powerful high-capacity, high-performance switch with wide geographical coverage is maintaining network efficiency in the face of constantly fluctuating traffic volumes. In response to this challenge, the Applicant also invented a self-configuring data switch comprising a number of electronic switching modules interconnected by a single-stage channel switch that includes a number parallel space switches, each having input ports and output ports. This switch architecture is described in Applicant's copending European Patent Application entitled SELF-CONFIGURING DISTRIBUTED SWITCH which was filed on April 5, 2000 and assigned Application No. 00302888.3. Each of the electronic modules is capable of switching variable-sized packets and is connected to the set of parallel space switches by a number of optical channels, each of the optical channels being a single wavelength in a multiple wavelength fiber link. The channel switching core permits any two modules to be connected by an integer number of channels. In order to enable the switching of traffic at arbitrary transfer rates, the inter-module connection pattern is changed in response to fluctuations in data traffic load. However, given the speed of optical switching equipment and the granularity of the channels, it is not always possible to adaptively modify the paths between modules to accommodate all data traffic variations. Consequently, it sometimes proves uneconomical to establish under-utilized paths for node pairs with low traffic volumes. To overcome this difficulty, a portion of the data traffic flowing between a source module and a sink module is switched through one or more intermediate nodes. Thus, in effect, the switch functions as a hybrid of a channel switch and linked buffer data switch, benefiting from the elastic path capacity of the channel switch.

[0006] A concentration of switching capacity in one location is, however, undesirable for reasons of security and economics. The self-configuring distributed switch with a high capacity optical core described in Applicant's co-pending Patent Application is limited in capacity and limited to switching entire channels. Consequently, it is desirable to provide a high-capacity switch with a distributed core. Such a core has the advantages of being less vulnerable to destruction in the event of a natural

disaster, for example. It is also more economical because strategic placement of distributed core modules reduces link lengths, requires less concentration of infrastructure and provides shorter paths for localized data traffic.

[0007] There therefore exists a need for a very high-capacity packet switch with a distributed core that is adapted to provide grade of service and quality of service guarantees. There also exists a need for a very high-capacity packet switch that provides intra-switch data paths of a finer granularity to reduce or eliminate a requirement for tandem switching.

SUMMARY OF THE INVENTION

[0008] According to the invention, there is provided a high capacity packet switch comprising a plurality of core modules, each of the core modules operating in a circuit switching mode, a plurality of edge modules connected to subtending packet sources and subtending packet sinks, each of the edge modules operating in a data packet switching mode. **CHARACTERIZED** by:

the core modules switch the data packets between the edge modules using wavelength division multiplexing (WDM) and time division multiplexing (TDM).

[0009] The high-capacity packet switch of the invention switches data packets between the edge modules using wavelength division multiplexing (WDM) and time division multiplexing (TDM). The switch has a distributed core that is adapted to provide guaranteed grade of service and quality of service. The switch of the invention can support intra-switch data paths with a granularity that reduces or eliminates a requirement for tandem switching.

[0010] The plurality of core modules can be geographically distributed and can operate independently. They may also periodically reconfigure to adapt to fluctuations in the network traffic. The core modules may coordinate reconfiguration with the edge modules in order to minimize a reconfiguration guard time.

[0011] Each of the core modules is preferably a space switch. Any of the well known textbook designs for a space switch can be used. However, the preferred space switch is an electronic single-stage rotator switch, because of its simple architecture, ease of control and scalability. The core modules preferably have neither input nor output buffers. A one of the edge modules is preferably co-located with each core module and serves as a controller for the core module.

[0012] Each of the edge modules has a plurality of ingress ports and a plurality of egress ports. Each of the ingress ports has an associated ingress queue. An ingress scheduler sorts packets arriving in the ingress queues from the subtending packet sources, the sort being by destination edge module from which the respective packets are to egress from the high capacity packet switch for delivery to the subtending packet sinks. The ingress scheduler periodically determines a number of packets waiting in the ingress queues for each other respective edge module, and sends a payload traffic allocation vector to each of the controllers of the core modules. The traffic allocation vector sent to a given controller relates only to a group of channels that connect the edge module to the given core module.

[0013] Each edge module also maintains a vector of pointers to the sorted payload packets, the vector of pointers being arranged in egress edge module order. A scheduling matrix for each slot in a time frame and each egress edge module is associated with the vector of pointers and determines a data transfer schedule for the ingress edge module.

[0014] Each ingress edge module also maintains an array of reconfiguration timing circuits, a one of the reconfiguration timing circuits being associated with each of the core modules. The reconfiguration timing circuits are respectively synchronized with time clocks in the respective edge modules that serve as controllers for the core modules. The reconfiguration timing circuits enable reconfiguration of channel switching in the core modules using a short guard time.

[0015] Each core module preferably comprises a plurality of rotator switches. Each rotator switch preferably accommodates a number of input channels equal to the number of edge modules, as well as a number of output channels equal to the number of edge modules. In a folded edge module configuration, each edge module preferably has one channel connected to an input port and one channel connected to an output port of each rotator switch. In an unfolded edge module configuration, each edge module is either an ingress module or an egress module. The ingress and egress modules are preferably arranged in co-located pairs. In the unfolded configuration, each ingress edge module preferably has one channel connected to an input port of each rotator switch. Each egress module likewise preferably has one channel connected to an output port of each rotator switch.

[0016] The invention also provides a method of switching payload data packets through a distributed data packet switch, comprising steps of receiving a payload data packet from a subtending source at an ingress edge module of the distributed data packets switch, determining an identity of an egress edge module from which the data packet should egress from the distributed data packet switch, arranging the data packet in a sorted order with other data packets received so that the data packets are arranged in a sorted order corresponding to the identity of the edge module from which the data packet should egress from the distributed data packet switch, transferring the sorted data packets in fixed-length data blocks to a core module of the distributed data packet switch, switching the fixed-length data blocks through the core module to the egress module, reconstructing the data packet at the egress edge module, and transferring the data packet from the egress module to a subtending sink **CHARACTERIZED** by:

using wave division multiplexing (WDM) and time division multiplexing (TDM) to provide intra-switch data paths of a fine granularity to reduce an requirement for tandem switching in the distributed data packet switch.

[0017] An ingress scheduler may periodically determine a number of packets waiting in ingress queues for each respective egress edge module and sends a capacity request vector to each of the controllers of the core modules, the capacity request vector sent to a given controller relating only to a group of channels that connect the edge module to the given core module.

[0018] Each ingress edge module may maintain a vector of pointers to the packets sorted by egress edge module and a scheduling matrix that provides a port number for each slot in which a data block can be transferred, the scheduling matrix being arranged in the same egress edge module order so that the scheduling matrix and the pointers are logically aligned; and, when a non-blank entry in the scheduling matrix indicates an egress port through which a data block can be transferred, a corresponding pointer in the vector of pointers is used to locate a starting point for the data block in the packets waiting in the ingress queues.

[0019] Preferably, an array stores pointers to packets sorted by egress edge module, and the pointers are dynamically updated each time a data block is transferred from the ingress queues to an egress channel so that reconfigurations of the core modules to the data can be accomplished without queue flushing. Two scheduling matrices, one in current use and one in update mode, are maintained at each ingress module. Each time a core reconfiguration occurs, a scheduling matrix in use is swapped for a current scheduling matrix. An unused copy of the scheduling matrix is available for update after the core reconfiguration. Rows in the matrix are executed sequentially, one per time slot, until a next core module reconfiguration occurs. After core module reconfiguration, processing continues at a next time slot.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] The invention will now be explained by way of example only, and with reference to the following drawings, in which:

FIG. 1 is a schematic diagram of a high capacity WDM-TDM packet switch in accordance with the invention having a centralized core;

FIG. 2 is a schematic diagram of the high capacity WDM-TDM packet switch shown in FIG. 1 wherein the space switches in the core are single-stage rotator switches;

FIG. 3 is a schematic diagram of a high capacity WDM-TDM packet switch in accordance with the invention with a distributed core;

FIG. 4 is a schematic diagram of a high capacity WDM-TDM packet switch in accordance with the invention showing an exemplary distribution of the core modules and edge modules;

FIG. 5 is a schematic diagram of a data structure used in each edge module to facilitate a process of computing capacity request vectors in the edge modules;

FIG. 6 is a schematic diagram of a table used by an ingress edge module to determine a preferred core module for a connection to an egress module;

FIG. 7 is a schematic diagram of data structures used in each core module for capacity scheduling using capacity request vectors received from the edge modules;

FIG. 8 is a schematic diagram illustrating space switch occupancy in a four core-module distributed switch in which a matching method employing a packing-search discipline is used; and

FIG. 9 is a schematic diagram of data structures used to control the transfer of data blocks from an ingress module to core modules of a high capacity WDM-TDM packet switch in accordance with the invention.

[0021] It should be noted that throughout the appended drawings, like features are identified by like reference numerals.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

[0022] FIG. 1 is a schematic diagram of a high-capacity WDM-TDM packet switch in accordance with the invention, generally indicated by reference 20. The packet switch 20 includes a plurality of edge modules 22, 24 shown for clarity of illustration in an "unfolded" configuration. In the unfolded configuration shown in FIG. 1, ingress edge modules 22 and egress edge modules 24 are separate switching modules constructed, for example, as described in Applicant's copending Patent Application Serial No. 00300700.2 which was filed January 31, 2000 and entitled RATE-CONTROLLED MULTICLASS HIGH-CAPACITY PACKET SWITCH, which describes the architecture and control of a high-capacity packet switch based on a rotator architecture. In a folded switch configuration, the ingress edge modules 22 and the egress edge modules 24 are combined into integrated switch modules of one ingress module and one egress module each, each integrated module having as many data ports as a sum of the data ports of the ingress edge module 22 and the egress edge module 24.

[0023] Located between the edge module pairs 22, 24 are a plurality of space switches 26 which serve as centralized core modules for the WDM-TDM packet switch 20. For the sake of scalability and switching speed, the space switches 26 are preferably electronic space switches, although optical space switches could be used and may become preferred when optical switching speeds improve. The space switches 26 are arranged

in parallel and, as will be described below, are preferably distributed in collocated groups. The number of edge modules 22, 24 and the number of space switches 26 included in the WDM-TDM packet switch 20 is dependent on the switching capacity required. In the example shown in FIG. 1, there are 256 (numbered as 0-255) ingress edge modules 22 and 256 egress edge modules 24. Each edge module 22 has egress ports to support 128 channels. In a typical WDM multiplexer, 16 wavelengths are supported on a link. Each wavelength constitutes a channel. Consequently, the 128 channels can be supported by eight optical fibers if cross-connectors are used, as will be explained below with reference to FIG. 3.

[0024] In order to ensure that any edge module 22 is enabled to send all of its payload traffic to any edge module 24, if so desired, each space switch 26 preferably supports one input channel for each module 22 and one output channel for each module 24. Therefore, in the example shown in FIG. 1, each space switch preferably supports 256 input channels and 256 output channels. The number of space switches 26 is preferably equal to the number of channels supported by each edge module 22, 24. In the example shown in FIG. 1, there are preferably 128 space switches 26, the number of space switches being equal to the number of channels from each ingress module 22.

[0025] FIG. 2 is a schematic diagram of a preferred embodiment of the WDM-TDM packet switch shown in FIG. 1. In accordance with a preferred embodiment, each of the space switches 26 is a single-stage rotator-based switch. In the rotator-based switch architecture, a space switch core is implemented as a bank of independent memories 28 that connect to the edge modules 22 of the switch through an ingress rotator 30. Traffic is transferred to the edge modules 24 of the switch 20 through an egress rotator 32. The two rotators 30, 32 are synchronized. A detailed description of the rotator switch architecture is provided in European Patent No. EP857383A1 that issued to Applicant on August 12, 1998, and provides a detailed description of the principles of operation of that switch architecture. In other respects, the switch architecture shown in FIG. 2 is identical to that shown in FIG. 1.

[0026] In the rotator switches 26, each bank of independent memories 28 is divided into a plurality of memory sections arranged in series. Each memory is preferably arranged in columns that are 128 bytes wide. Each memory is divided into a number of partitions, the number of partitions being equal to the number of egress edge module 24. The size of the memory portion governs a size of data block switched by the channel switching core. The size of the data block is a matter of design choice, but is preferably about 4-16 kilobits.

Partitioning the Core

[0027] The channel switching core is preferably partitioned into core modules and distributed for two principal reasons: economics and security. FIG. 3 is a schematic diagram of a preferred embodiment of a distributed WDM-TDM packet switch in accordance with the invention. A plurality of core modules 34 are geographically distributed. A plurality of cross-connectors 36, which may be, for example, very slow optical switches, connect a plurality of ingress and egress edge modules 22, 24 to the core modules 34. The cross-connectors 36 serve as multiplexers and thereby reduce the number of physical links required to connect each ingress and egress edge module 22, 24 to each core module 34. The core modules 34 preferably include an equal number of rotator switches. A WDM-TDM packet switch 20 of a size shown in FIGs. 1 and 2, with eight core modules 34, includes 16 rotator switches 28 in each core module 34 when geographically distributed as shown in FIG. 3 (8 x 16= 128). If the ingress and egress edge modules 22, 24 are grouped in clusters of eight per cross-connector 36, then 64 cross-connectors are required to connect the ingress and egress edge modules 22, 24 to the core modules 34. The clustering of the ingress and egress edge modules 22, 24 and the number of cross-connectors 36 used in any given installation is dependent on network design principles well understood in the art and does not require further explanation. In any distributed deployment of the WDM-TDM packet switches, it is preferred that each ingress and egress edge module 22, 24 be connected to each space switch 26 of each core module 34 by at least one channel. The switch may be partitioned and distributed as desired with the exception that one ingress and egress edge module 22 is preferably collocated with each core module 34 and serves or hosts a controller, as a controller for the core module, as will be explained below in more detail.

[0028] FIG. 4 shows an exemplary distribution of a WDM-TDM packet switch 20 in accordance with the invention, to illustrate a hypothetical geographical distribution of the switch. Cross-connectors 36 and optical links 38 are not shown in FIG. 4 for the sake of clarity. In this example, 16 ingress and egress edge modules 22, 24 numbered 0-15 and four core modules 34 numbered 0-3 are potentially distributed over a large geographical area. As explained above, an ingress edge module 22 is collocated with each core module 34. In this example, ingress edge modules 0-3 are collated with corresponding core modules 0-3. Because the space switches 26 are rote devices with substantially no computational capacity, they require controllers to perform scheduling allocations and other functions which are described below in more detail. The ingress edge modules 22 include high-speed processors which are capable of performing control functions, or hosting control functions, for the core modules 34. Consequently, an ingress edge module 22 is preferably collocated with each core module 34. The processor of the ingress edge module 22 need not, however, perform the control functions of the core module 34. Rather, it may host, at one

of its ports, a processor to perform the control functions of the core module 34. The collocation is also important to enable time coordination in the distributed WDM-TDM packet switch 20, as explained below.

Time Coordination in the Distributed WDM-TDM Packet Switch

[0029] Time coordination is required between ingress edge modules 22 and core modules 34 if the WDM-TDM packet switch 20 geographically distributed. Time coordination is necessary because of propagation delays between ingress edge modules 22 and the core modules 34. Time coordination is accomplished using a method described in Applicant's above-referenced copending patent application filed April 4, 1999. In accordance with that method, time coordination is accomplished using an exchange of timing packets between the ingress edge modules 22 and the respective edge module controller for core modules 34. At predetermined intervals, each ingress edge module 22 is programmed to send a timing packet to the ingress edge module 22 that serves as a controller for the associated core module 34. For example, ingress edge module 9 (FIG. 4) at a predetermined interval sends a timing packet to ingress edge module 3 associated with core module 3. On receipt of the timing packet, the ingress edge module 3, which serves as a controller for the core module 3, stamps the packet with a time stamp that indicates its local time. At some convenient time prior to the next predetermined interval, the time stamped packet is returned to the edge module 9. The edge module 9, and each of the other ingress edge modules 0-15, maintains an array of *M* reconfiguration timing circuits where *M* equals the number of core modules 34. The core modules 34 operate independently and reconfigure independently. as will be described below in more detail. Consequently, each ingress edge module 22 must maintain a separate reconfiguration timing circuit coordinated with a local time of an ingress edge module 22 collocated with each core module 34. Without timing coordination, guard times for reconfiguration of the core modules 34 would have to be too long due to the propagation delays between the geographically distributed ingress edge modules 22 and the core modules 34.

[0030] For example, in the configuration of the WDM-TDM packet switch 20 shown in FIG. 4, each ingress edge module 22 must maintain an array of four reconfiguration timing circuits respectively coordinated with the local times of ingress edge modules 0-3 collocated with the respective core modules 34. As explained above, in order to maintain time coordination, the ingress edge module 9. at regular predetermined intervals, sends a timing packet to the ingress edge module 0. The timing packet is sent over a communications time slot and received on an ingress port of the ingress edge module 0 dedicated to management functions. The ingress port, on receipt of the timing packet, time stamps the packet with the time from its local time (timing circuit 0) and queues the timing packet for return to the edge module 9. At some convenient later time before the start of the next timing interval, the timing packet is returned to the ingress edge module 9. On receipt of the timing packet at ingress edge module 9, the ingress edge module 9 uses the time at which the packet was received at ingress edge module 0 (time stamp) in order to coordinate its reconfiguration timing circuit 0 with the local time of ingress edge module 0. Several methods for timing coordination are explained in detail in Applicant's copending Patent Application Serial No. 09/286.431 filed April 6, 1999.

Packet Transfer Through the WDM-TDM Packet Switch

[0031] Ingress and egress edge modules 22, 24 of the WDM-TDM packet switch 20 operate in packet switching mode. The edge modules 22, 24 are adapted to switch variable sized packets and transfer the packets to subtending sinks in the format in which the packets were received. Switching in the core modules 34 is accomplished in circuit switching mode. The core modules 34 are completely unaware of the content switched and simply switch data blocks. In order to improve resource allocation granularity, the WDM-TDM packet switch 20 switches in both wave division multiplexing (WDM) and time division multiplexing (TDM) modes. Each link 38 (FIG. 3) interconnecting the switched edge modules 22, 24 and the core modules 34 is preferably an optical link carrying WDM data on a number of channels, each channel being one wave length in the WDM optical link 38. Each channel is further divided into a plurality of discrete time slots, hereinafter referred to simply as "slots". The number of slots in a channel is a matter of design choice. In a preferred embodiment, each channel is divided into 16 time slots. Consequently, the smallest assignable increment of bandwidth is $1/16^{th}$ of the channel capacity. For a 10 gigabit per second (10 Gb/s) channel, the smallest assignable capacity allocation is about 625 megabits per second (625 Mb/s). Connections requiring less capacity are aggregated by class-of-service and quality-of-service in a manner well known in the art. Connections requiring more capacity are allocated multiple slots, as required.

Admission Control

[0032] The capacity requirement for each connection established through the WDM-TDM packet switch 20 is determined either by a specification received from a subtending source or, preferably, by automated traffic measuring mechanisms based on traffic monitoring and inference. If automated measurement is used, the capacity requirements are expressed as a number of slots for high bandwidth connections. For aggregated traffic, the capacity requirements are measured for a class of

service. Regardless of the method used to estimate the capacity requirements, it is the responsibility of the ingress edge modules 22 to quantify the traffic requirements for its traffic load. It is also the responsibility of the ingress edge modules 22 to select a route for each admission request. Route selection is accomplished using connection tables provided by a Network Management System (NMS) (not illustrated) which provides a table of preferred connecting core modules between each ingress and egress edge module.

[0033] Admission control may be implemented in a number of ways that are well known in the art, but the concentration of responsibility is at the edge and any ingress edge module 22 receiving an admission request first determines whether free capacity is available on any of the preferred routes through a core module defined in its connection table prior to acceptance.

Scheduling at the Edge

[0034] At any given time, each ingress edge module 22 has an allocated capacity to each egress edge module 24 expressed as a number of slots. The number of allocated slots depends on a capacity allocation, which may be 0 for certain ingress/egress module pairs. The allocated capacities may be modified at regular reconfiguration intervals which are independently controlled by the controllers of the distributed core modules 34. An ingress edge module 22 accepts new connections based on its current capacity allocation to each egress edge module 24. The controller of each ingress edge module 22 also monitors its ingress queues, which are sorted by egress edge module, as described above, to determine whether a change in capacity allocation is warranted. It is the responsibility of each ingress edge module 22 to determine when resources should be allocated and when resources should be released. However, it is the controllers at the core modules 34 that determine whether a bandwidth allocation request can be granted. Bandwidth release requests are always accepted by the controllers of the core modules 34. The re-allocation of bandwidth and the reconfiguration of the core modules 34 is explained below in more detail.

[0035] Each ingress edge module 22 determines its bandwidth capacity requirements and communicates those requirements to the controllers of the respective core modules 34. On receipt of a capacity requirement, a controller of a core module 34 attempts to satisfy the requirement using a rate matching process. The controller of the core modules 34 compute a scheduling matrix based on the capacity requirements reported by each ingress edge module 22, as will be explained below, and returns relevant portions of the scheduling matrix to each ingress edge module 24 prior to a reconfiguration of the core module 34. At reconfiguration, three functions are implemented. Those functions are: a) releases, which return unused resources to a resource pool; b) bandwidth capacity increases which allocate new bandwidth to ingress edge modules 22 requiring it; and c) new bandwidth capacity allocations, in which an allocation for an ingress edge module 22 is increased from 0.

Capacity Scheduling

[0036] As described above, the ingress edge modules 22 are responsible for managing their bandwidth capacity requirements. Consequently, each edge module computes a capacity requirement vector at predetermined intervals such that the capacity requirement is reported to each core module 34 at least once between each core reconfiguration. FIG. 5 illustrates the computation of the capacity requirement vector. As shown, an ingress edge module 22 constructs a matrix of x rows and y columns, where x is the number of ingress ports and y is the number of egress modules in the switch configuration. In the example shown, the number of channels is 128 and number of ingress edge modules 22 is 255. A number representative of an actual count of packets in the egress buffers, or a number resulting from a traffic prediction algorithm, is inserted in each cell of the matrix shown in FIG. 5. A traffic allocation requirement sum 40 provides a summation for each egress edge module 24 of the total capacity requirement. The total capacity requirement is then subdivided into M capacity requirement vectors, where M is the number of core modules 34 and the respective capacity requirement vectors are distributed to the respective core modules to communicate the capacity requirements. A zero in a capacity requirement vector indicates that any bandwidth capacity previously allocated to the ingress core module 22 is to be released.

[0037] In order for an ingress edge module 22 to intelligently request a bandwidth capacity increase, it must follow a governing procedure. As described above, each ingress edge module 22 is provided with a table of preferred connections to each egress edge module 24. FIG. 6 shows how the table of preferred connections through the switch is used in the bandwidth allocation request process. A preferred connection table 42 is provided to edge module 7 in the network shown in FIG. 4. The preferred connection table 42 provides the edge module 7 with the core modules through which connections can be made with egress edge modules, the core module numbers being listed in a preferred order from top to bottom. Each entry 44 in the preferred connection table 42 is a core module identification number. Therefore, if ingress edge module 7 needs to send packets to egress edge module 0, the preferred core module for the connection is core module 0. The other core modules that may be used for the connection are, in descending order of preference, 3, 1 and 2. Likewise, if edge module 7 needs to send packets to edge module 15, the preferred core module is core module 3, and the alternate core modules, in descending preference, are 2, 0 and 1.

[0038] As shown in FIG. 6, the preferred connection table 42 is used in each edge module to facilitate the process of requesting bandwidth capacity allocations from the respective core modules 34. The array 40 of the traffic allocation summary computed as described above, has 16 entries, one entry for traffic destined to each egress edge module. The array is matched with the preferred connection table 42, which has 16 columns and four rows, as explained above. The array 40 indicates the number of slots required to accommodate traffic from the edge module 7 to the 15 other edge modules in the network shown in FIG. 4. These data structures are used to construct the capacity request vectors described above, which are sent to the respective core modules 34. As will be explained below in more detail, reconfiguration of the core modules is preferably staggered so two core modules do not reconfigure at the same time. Consequently, there is a staggered reconfiguration of the core modules 34. For each capacity request vector sent by an ingress edge module 22, a first set of capacity request vectors is preferably constructed using the preferred connections listed in row 1 of the preferred connection table 42. If a capacity request denial is received back from a core module, an updated capacity request vector is sent to a second choice module. In planning capacity allocations prior to reconfiguration, a core module preferably uses the last received allocation request vector until processing has advanced to a point that any new capacity request vectors are ignored. Consequently, for example, the capacity request vector sent to core module 0 would request five slots for a connection to egress edge module 0, seven slots for a connection to edge module 11, seven slots for a connection to edge module 13, and ten slots for a connection to edge module 14. If core module 0 denied any one of the capacity requests, an updated capacity request vector would be sent to the next preferred core module shown in the preferred connection table 42.

[0039] FIG. 7 illustrates a scheduling function performed by each of the controllers for the respective core modules 34. Each controller for the respective core modules 34 receives capacity request vectors from the ingress edge modules 22. The capacity request vectors received from each ingress edge module 22 is expressed in terms of the number of slots that each ingress edge module requires to accommodate its traffic switched through the given core module 34. The controller of each core module 34 assembles the capacity request vectors in a traffic allocation request matrix 44 which includes N rows and N columns where N equals the number of ingress edge modules. In the example network shown in FIG. 4, the traffic allocation request matrix 44 constructed by the controller of each core module 34 would be a 16 x 16 matrix (256 x 256 matrix for the network shown in FIG. 3).

[0040] The traffic allocation request matrix 44 is normally a sparse matrix with a majority of null entries. The controller for the core module attempts to schedule the capacity requested by each ingress edge module 22 using data structures generally indicated by references 46 and 48. Each of the data structures 46, 48 is a three-dimensional matrix having a first space dimension s, which represents the respective space switches associated with the core module 34; a second space dimension p, which represents the input channels; and a time dimension t, which represents the slots in a slotted frame. Thus, an entry in data structure 46 is represented as $\{s,p,t\}$. The second dimension p may represent an input channel, if associated with the data structure 46, or an output channel if associated with the data structure 48. If the number of slots T per frame is 16, for example, then in the configuration of FIG. 1, which shows a centralized core, the size the three-dimensional structure 46 is 256 x 128 x 16. In the distributed core shown in FIG. 3, each core module uses a three-dimensional structure 46 of size 256 x 16 x 16.

[0041] When the connections through a core module 34 are reconfigured, the core controller may either reschedule the entire capacity of the respective core module 34 or schedule bandwidth capacity changes by simply perturbating a current schedule. If the entire bandwidth capacity of the core module is reconfigured, each ingress edge module 22 must communicate a complete capacity request vector to the core module while, in the latter case, each ingress edge module 22 need only report capacity request changes, whether positive or negative, to a respective core controller. A negative change represents capacity release while a positive change indicates a request for additional bandwidth capacity. The incremental change method reduces the number of steps required to prepare for reconfiguration. However, the incremental change method potentially risks allocation efficiency, because a core module that fails to accommodate a bandwidth capacity request may force the requesting ingress edge module to seek incremental capacity allocation from another, less than optimal, connection through another core module that may lead to a longer path to a destination.

[0042] The capacity scheduling done by the controller for a core module 44 can be implemented by simply processing the non-zero entries in the traffic allocation request matrix 44 one at a time. A non-zero entry 50 in the traffic allocation request matrix 44 represents a number of required slots for a respective edge module pair. A three dimensional data structure 46 indicates free input slots at core modules, and data structure 48 shows the free slots at output ports of the core module 34. The three dimensional data structures 46, 48, initialized with null entries, are then examined to determine if there are sufficient matching slots to satisfy the capacity request. Each cell 51 in each data structures 46, 48 represents one slot. A slot in structure 46 and a slot in structure 48 are matching slots if each is unassigned and if both have the same first space dimensions (s) and time dimension (t). Thus, slot $\{s,j,t\}$ in data structure 46 and slot $\{s,k,t\}$ in data structure 48 are matching if both are free, with

respect to the values of *j* and *k*.

[0043] A bandwidth capacity request is rejected by a core module if sufficient matching slots cannot be found. In order to reduce the incidence of mismatch, the matching process should always start from a selected space switch at a selected time slot and follow the same search path for each capacity request. For example, the matching process may start from space switch 0 at slot 0 and then proceed by incrementing the slot from 0 to *S*, where *S* is the number of time slots per channel. It then continues to the next time port plane 53 until the 16 planes (in this example) are exhausted or the capacity is successfully allocated, whichever takes place first. The result produced by this packing search, which is well known in the art, is an occupancy pattern shown in FIG. 8.

[0044] FIG. 8 shows a typical space switch occupancy for each of the core modules 34. Each core module 34 includes four space switches in this example. Observing any of the core modules, the occupancy of the space switch at which the matching search always starts is higher than the occupancy of the second space switch in the search path, etc. This decreasing occupancy pattern is known to provide a superior matching performance over methods that tend to equalize the occupancy, such as a round-robin or random search.

Packet Transfer from the Edge Modules to the Core

[0045] As a result of the scheduling process described above, each core module, prior to reconfiguration, returns to each ingress edge module 22 a schedule vector which is used to populate a schedule matrix 54 partially shown in FIG. 9. The schedule matrix 54 is a matrix containing *S* rows (where *S* = 16 in this example) and *N* columns where *N* equals the number of ingress edge modules 22. The 16 rows, only four of which are illustrated, represent the 16 slots in a frame. The non-blank entries 56 in the schedule matrix represent channel numbers of the egress channels of an egress edge module 22. The edge module is enabled to transfer one data block to a core module 34 for each valid entry in the schedule matrix 54. For example, in row 1 of the matrix 54 shown in FIG. 8, the ingress edge module 22 can transfer a data block through ports 16 to egress edge module 14. In time slot 2, the edge module can transfer one data block through channel 97 to edge module 1, and one data block through channel 22 to edge module 14. The ingress edge module 22 has no knowledge of the core module to which the data block is to be transferred and requires none.

[0046] The size of a data block is a matter of design choice, but in the rotator-based core modules, the size of a data block is related to the structure of middle memories 28 (FIG. 2). In general, a data block is preferably 4 kilobits (Kb) to about 16 Kb. In order for data blocks to be transferred from the ingress queues to the appropriate egress channel, an array 58 stores pointers to packets sorted by egress edge module. The pointers 58 are dynamically updated each time a data block is transferred from the ingress queues to an egress channel. Consequently, reconfigurations of the core modules 30 can be accomplished without queue flushing.

[0047] In actual implementations, it is preferable to maintain two matrices 54, one in current use and one in update mode. Each time a core reconfiguration takes place, the matrix in use is swapped for a current matrix. The unused copy of the matrix is available for update. Rows in the matrix are executed sequentially one per slot until a next core module reconfiguration occurs. After core module reconfiguration, processing continues at the next slot.

[0048] The invention thereby provides a very high-speed packet switch capable of wide geographical distribution and edge-to-edge total switching capacity that is scalable to about 320 Tera bits per second (Tbs) using available electronic and optical components. The control is principally edge-based and the core modules 34 operate independently so that if one core module fails, the balance of the switch continues to operate and traffic is rerouted through the remaining available core modules. Normal failover techniques well known in the art may be used to ensure continuous operation in the event of component failure.

[0049] The embodiments of the invention described above are intended to be exemplary only. The scope of the invention is therefore intended to be limited solely by the scope of the appended claims.

Claims

1. A high capacity packet switch comprising a plurality of core modules, each of the core modules (26) operating in a circuit switching mode, a plurality of edge modules (22, 24) connected to subtending packet sources and subtending packet sinks, each of the edge modules operating in a data packet switching mode, **CHARACTERIZED** by:

   the core modules switch the data packets between the edge modules using wavelength division multiplexing (WDM) and time division multiplexing (TDM).

2. A high capacity packet switch as claimed in claim 1 wherein each core module (26) is a space switch.

3. A high capacity packet switch as claimed in claim 2 wherein each core module (26) is a single-stage electronic rotator switch.

4. A high capacity packet switch as claimed in any preceding claim wherein:

   each edge module (22,24) has a plurality of ingress ports, each of the ingress ports having an

associated ingress queue; and
an ingress scheduler sorts packets arriving in the ingress queues from the subtending packet sources, the sort being by egress edge module (24) from which the respective packets are to egress from the high capacity packet switch (20) for delivery to the subtending packet sinks.

5. A high capacity packet switch as claimed in claim 4 wherein the ingress scheduler periodically determines a number of packets waiting in the ingress queues for each respective egress edge module (22,24) and sends a capacity request vector to each of the controllers (22) of the core modules (34), the capacity request vector sent to a given controller relating only to a group of channels that connect the edge module to the given core module.

6. A high capacity packet switch as claimed in claims 4 or 5 wherein each ingress edge module (22) maintains a vector of pointers (58) to the packets sorted by egress edge module (24) and a scheduling matrix (54) that provides a port number (56) for each slot in which a data block can be transferred, the scheduling matrix being arranged in the same egress edge module order so that the scheduling matrix and the pointers are logically aligned; and, when a non-blank entry in the scheduling matrix indicates an egress port through which a data block can be transferred, a corresponding pointer in the vector of pointers is used to locate a starting point for the data block in the packets waiting in the ingress queues.

7. A high capacity packet switch as claimed in any preceding claim wherein the core modules (34) and the edge modules (22,24) are spatially distributed.

8. A high capacity packet switch as claimed in claim 7 wherein one edge module (22) is co-located with each core module (34), and the edge module serves as a controller for the core module.

9. A high capacity packet switch as claimed in claims 7 or 8 wherein each edge module (22,24) has *M* reconfiguration timing circuits, where *M* is the number of core modules (34), each of the reconfiguration timing circuits being time-coordinated with a time counter in the respective edge modules that serve as processors for the core modules, to coordinate data transfer from the ingress edge modules when the core modules are reconfigured to change channel connectivity.

10. A high capacity packet switch as claimed in any preceding claim wherein each edge module (24,34) is connected to each core module by at least one communications link (25).

11. A high capacity packet switch as claimed in claim 10 wherein each core module (34) comprises a plurality of single-stage rotator switches (26), each rotator switch having a number of input ports collectively adapted to accommodate a number of channels equal to the number of ingress edge modules (22,24) and a number of output ports collectively adapted to accommodate a number of channels equal to the number of egress edge modules, and each edge module has at least one channel to each of the rotator switches.

12. A high capacity distributed packet switch as claimed in any preceding claim wherein the ingress edge module (22) and the egress edge modules (24) comprise integrated units of one ingress edge module and one egress edge module each.

13. A method of switching payload data packets through a distributed data packet switch (20), comprising steps of receiving a payload data packet from a subtending source at an ingress edge module (22) of the distributed data packets switch, determining an identity of an egress edge module (24) from which the data packet should egress from the distributed data packet switch, arranging the data packet in a sorted order with other data packets received so that the data packets are arranged in a sorted order corresponding to the identity of the edge module from which the data packet should egress from the distributed data packet switch, transferring the sorted data packets in fixed-length data blocks to a core module (34) of the distributed data packet switch, switching the fixed-length data blocks through the core module to the egress module, reconstructing the data packet at the egress edge module, and transferring the data packet from the egress module to a subtending sink **CHARACTERIZED** by:
using wave division multiplexing (WDM) and time division multiplexing (TDM) to provide intra-switch data paths of a fine granularity to reduce a requirement for tandem switching in the distributed data packet switch.

14. A method as claimed in claim 13 wherein an ingress scheduler periodically determines a number of packets waiting in ingress queues for each respective egress edge module (22) and sends a capacity request vector to each of the controllers (22) of the core modules (34), the capacity request vector sent to a given controller relating only to a group of channels that connect the edge module to the given core module.

15. A method as claimed in claim 14 wherein each ingress edge module (22) maintains a vector of pointers to the packets sorted by egress edge module

(24) and a scheduling matrix (54) that provides a port number for each slot in which a data block can be transferred, the scheduling matrix being arranged in the same egress edge module order so that the scheduling matrix and the pointers are logically aligned; and, when a non-blank entry in the scheduling matrix indicates an egress port through which a data block can be transferred, a corresponding pointer in the vector of pointers is used to locate a starting point for the data block in the packets waiting in the ingress queues.

16. The method as claimed in claim 15 wherein an array (58) stores pointers to packets sorted by egress edge module, and the pointers are dynamically updated each time a data block is transferred from the ingress queues to an egress channel so that reconfigurations of the core modules (30) can be accomplished without queue flushing.

17. The method as claimed in claims 15 or 16 wherein two scheduling matrices (54). one in current use and one in update mode, are maintained at each ingress module (22) .

18. The method as claimed in claim 17 wherein each time a core reconfiguration occurs, a scheduling matrix (54) in use is swapped for a current scheduling matrix.

19. The method as claimed in claim 17 or 18 wherein an unused copy of the scheduling matrix (54) is available for update after the core reconfiguration.

20. The method as claimed in claim 19 wherein rows in the matrix are executed sequentially, one per time slot, until a next core module reconfiguration occurs; and, after core module (34) reconfiguration, processing continues at a next time slot.

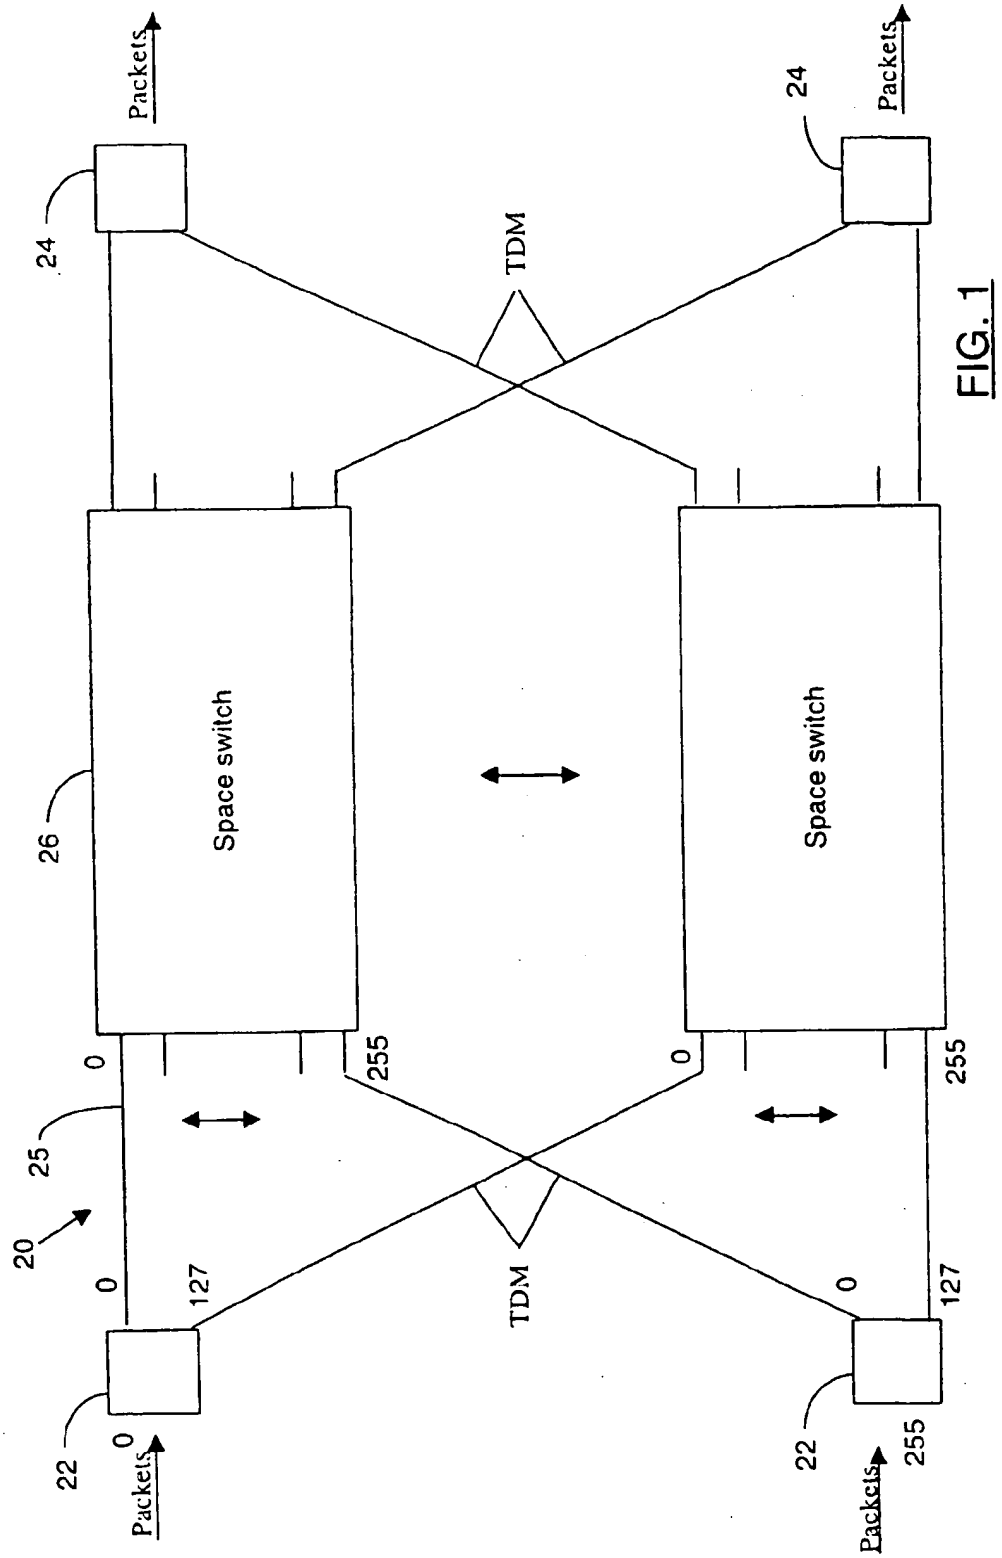

FIG. 1

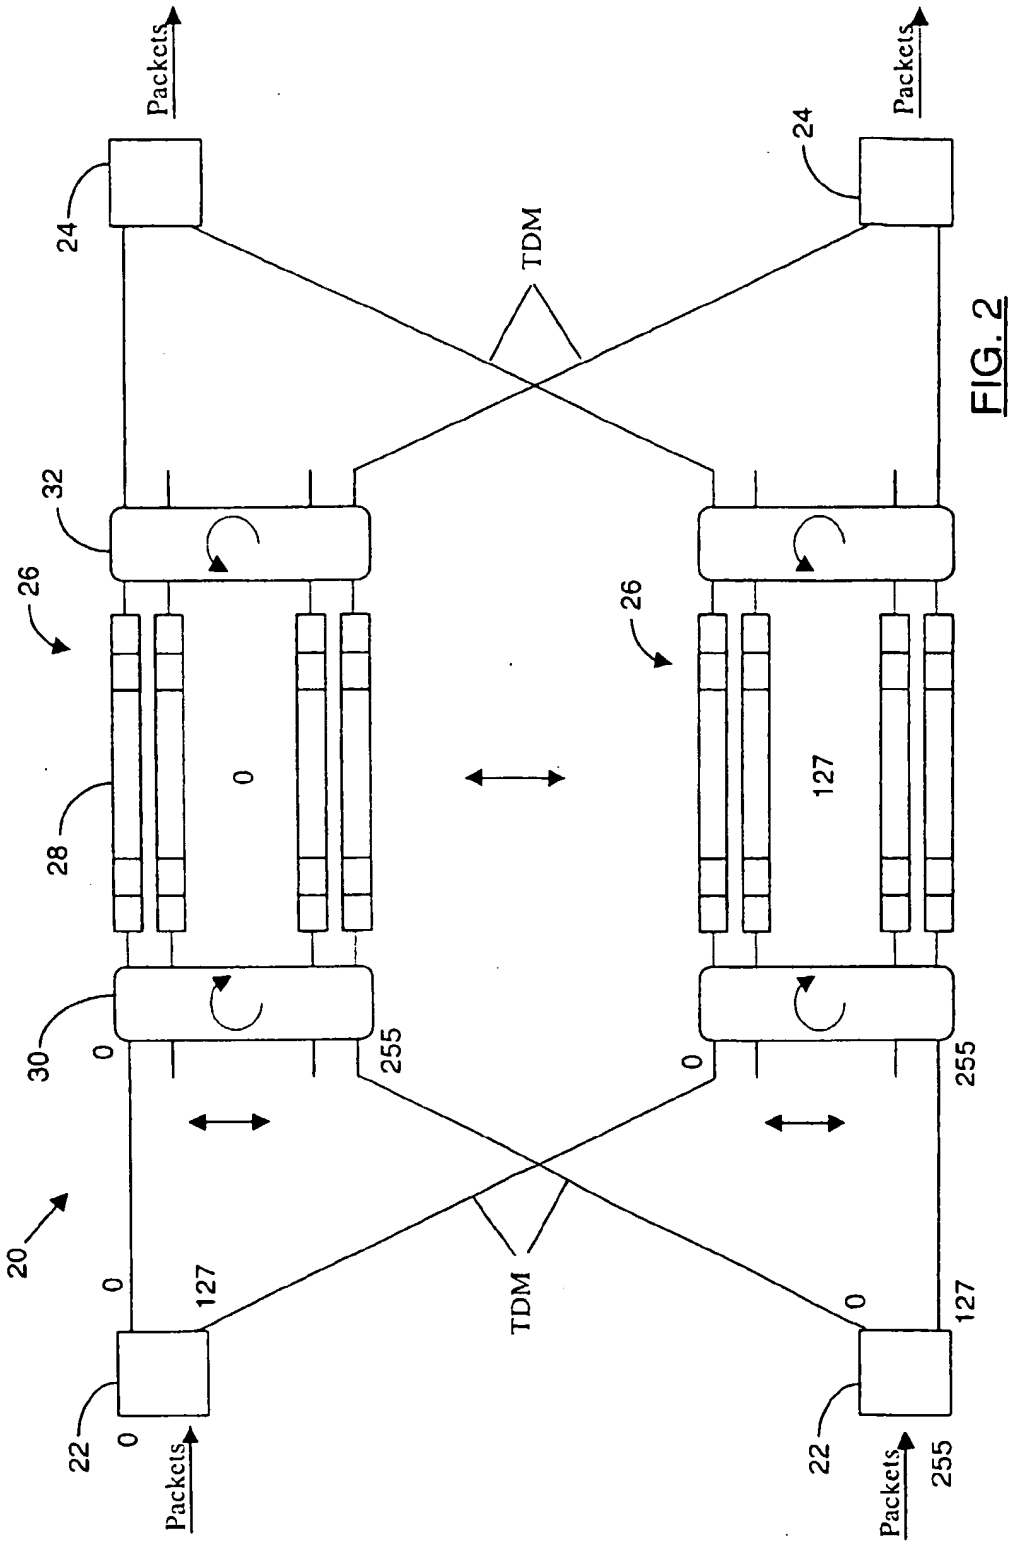

FIG. 2

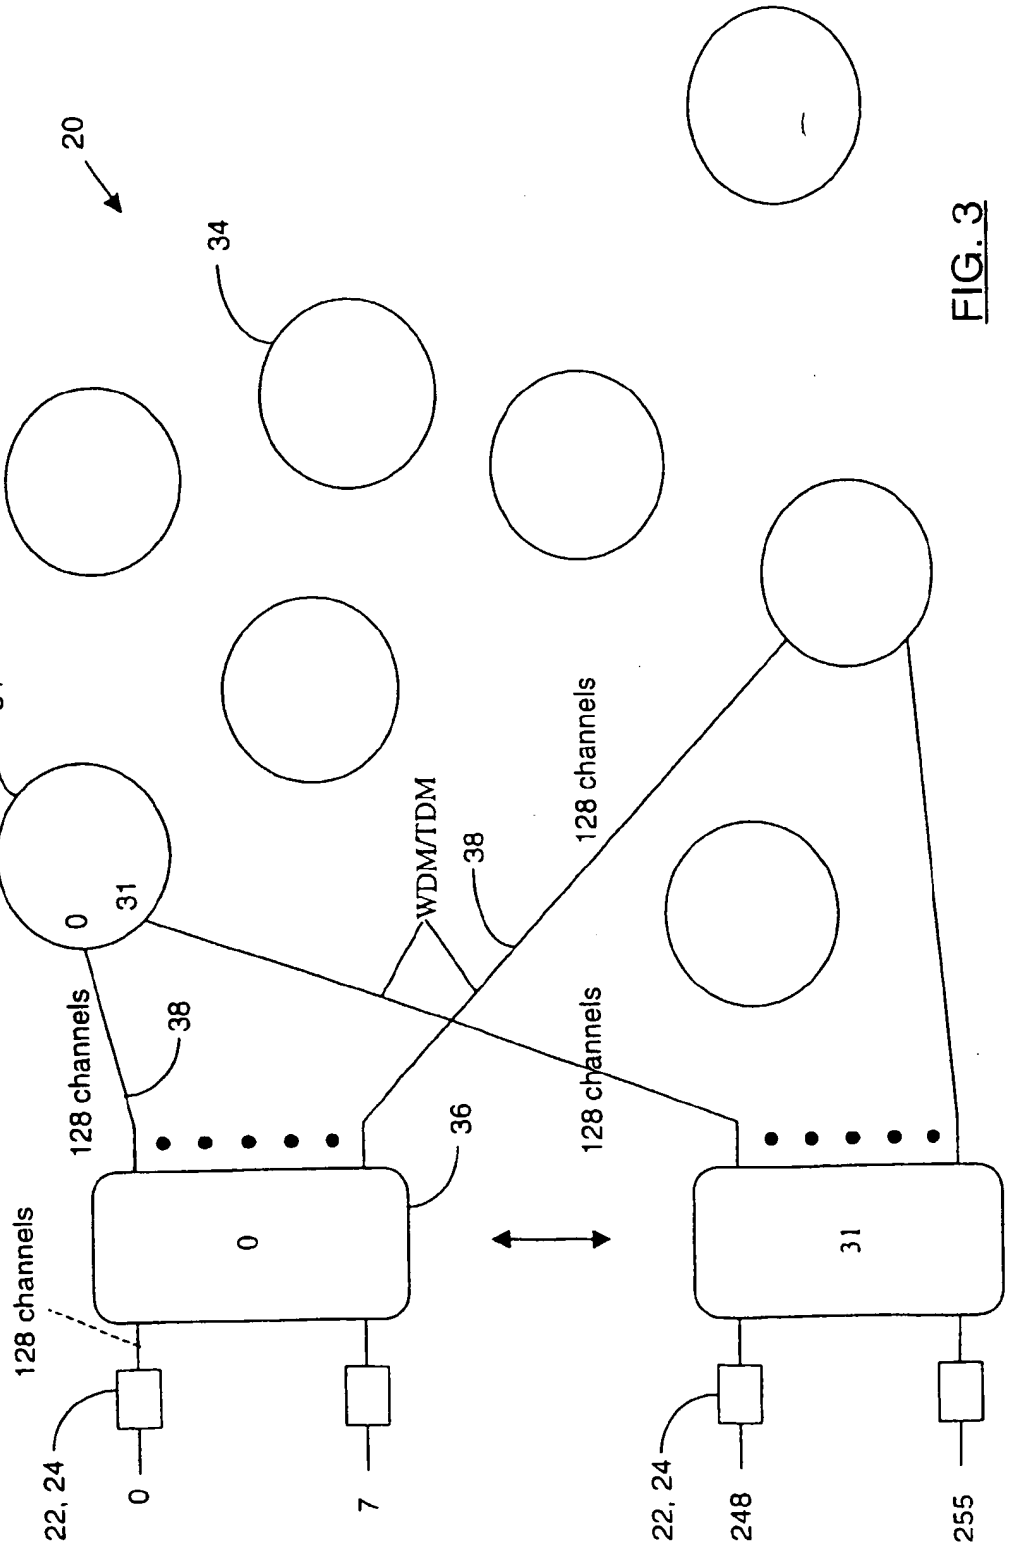
20
34
34
0
31
WDM/TDM
38
128 channels
128 channels
38
128 channels
36
128 channels
0
31
22, 24
0
7
22, 24
248
255

FIG. 3

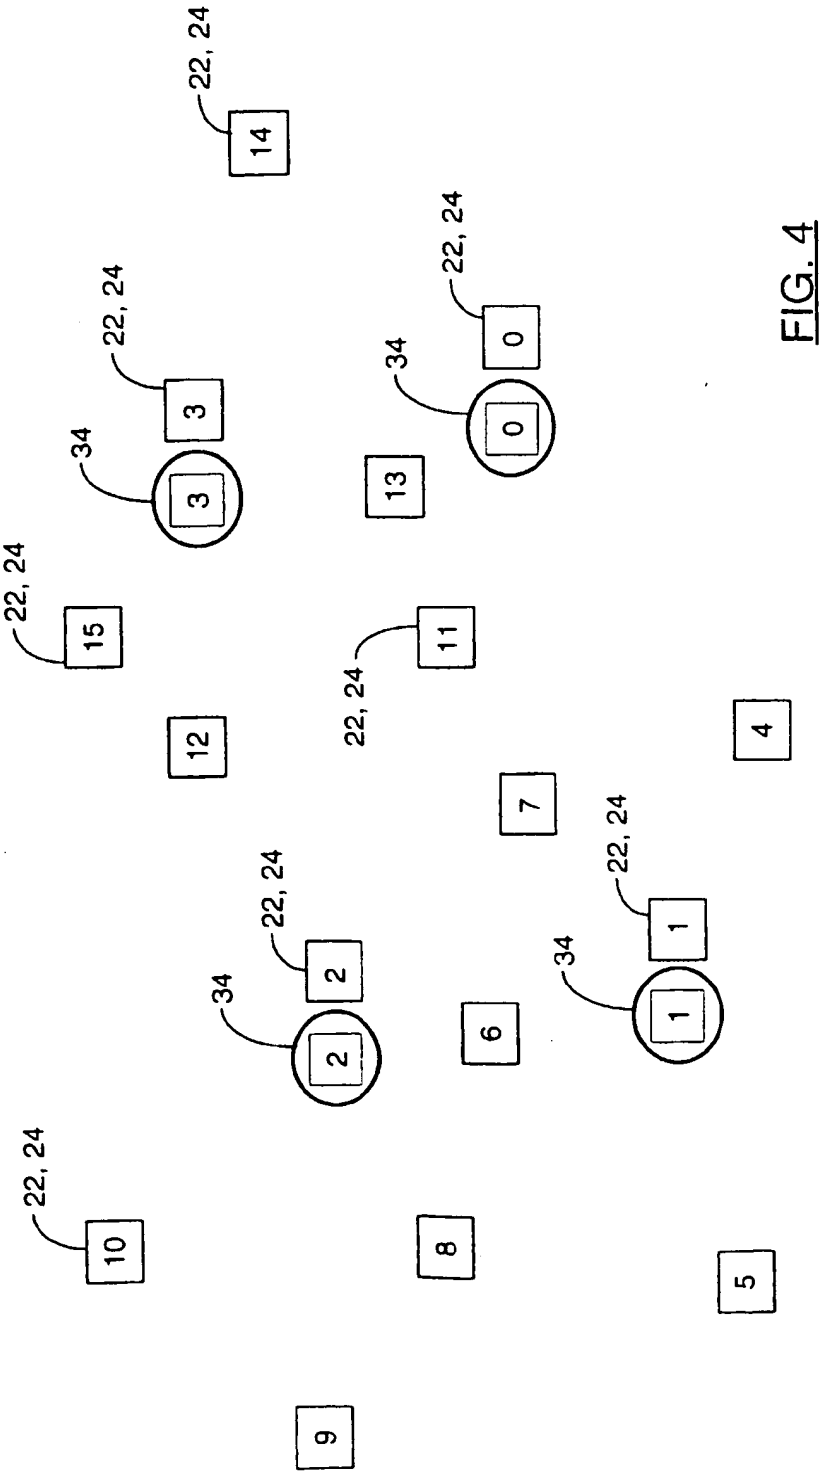

FIG. 4

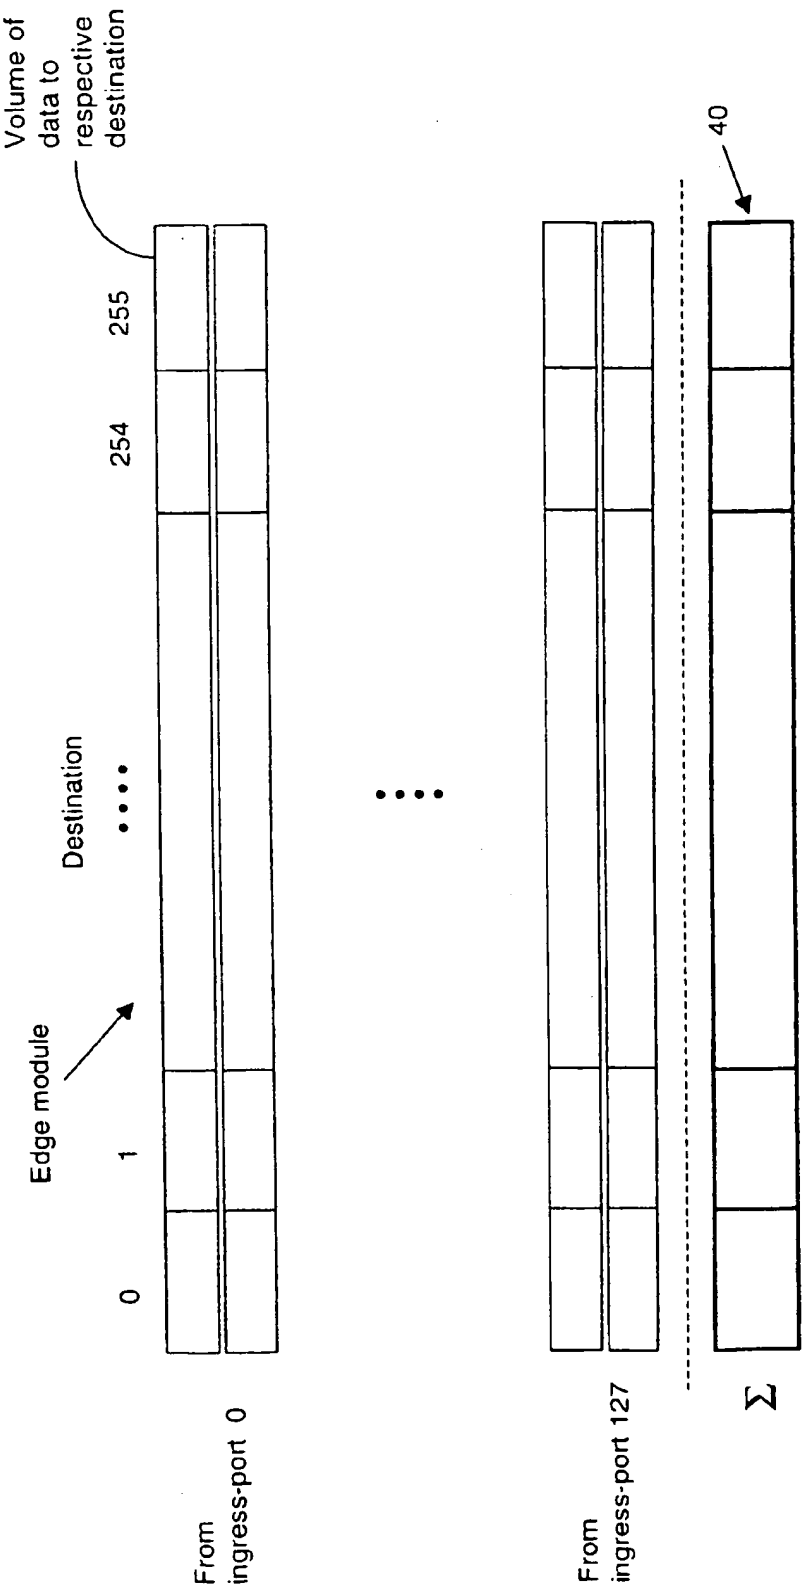

FIG. 5

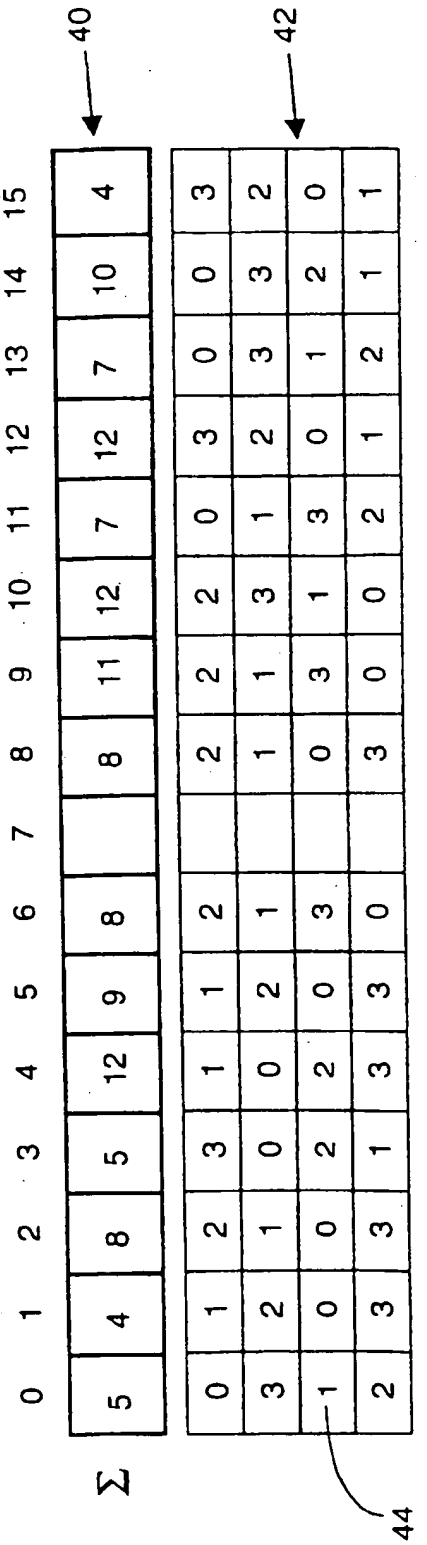

FIG. 6

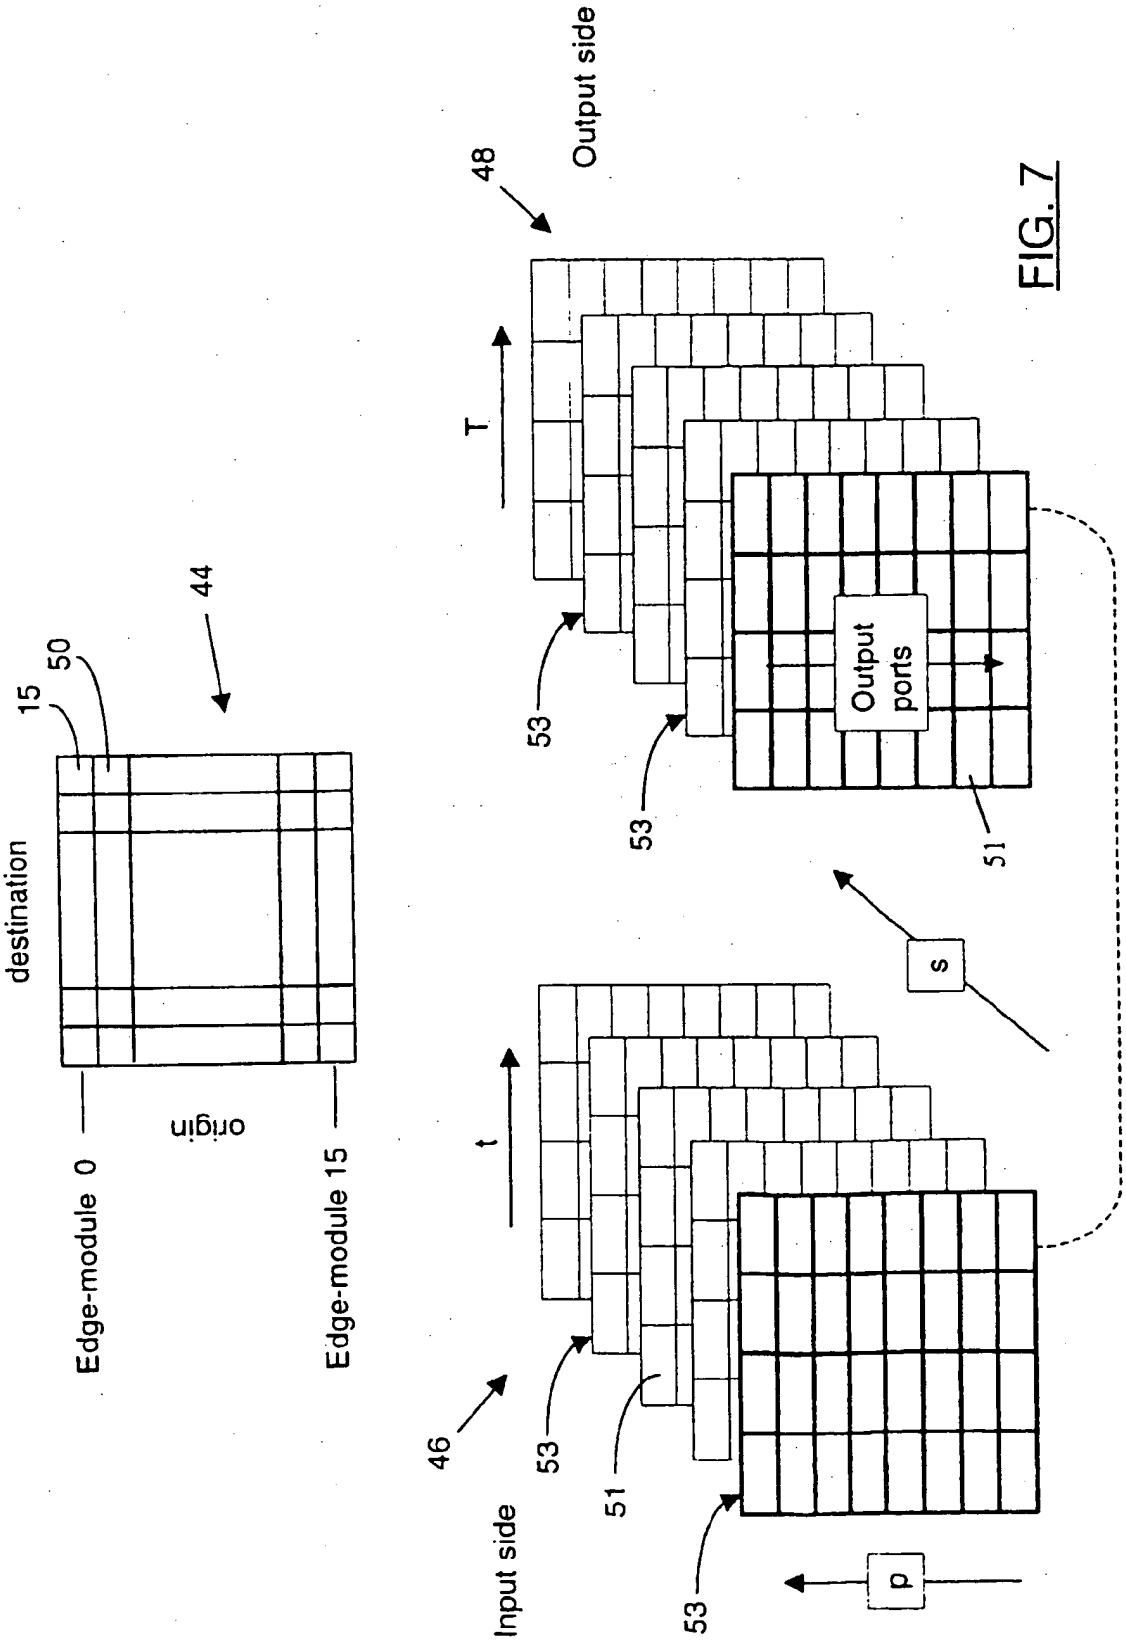
Output side
48
T
Output ports
53
53
51
s
t
46
53
51
Input side
53
p
44
15
50
destination
Edge-module 0
origin
Edge-module 15

FIG. 7

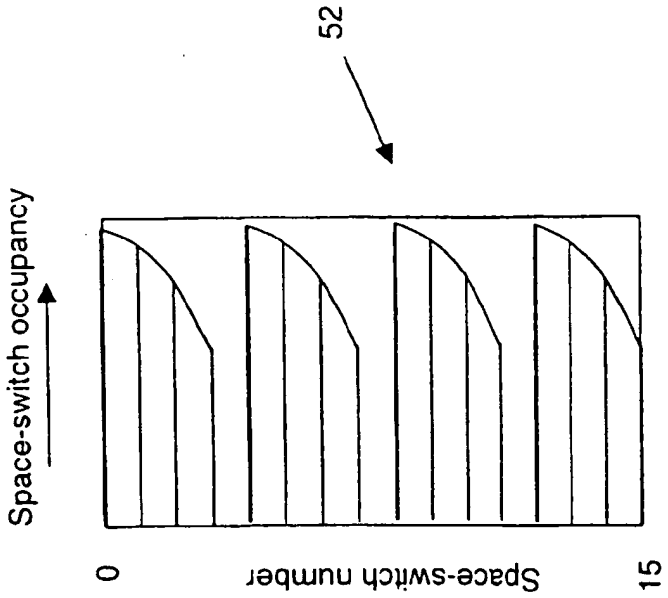

FIG. 8

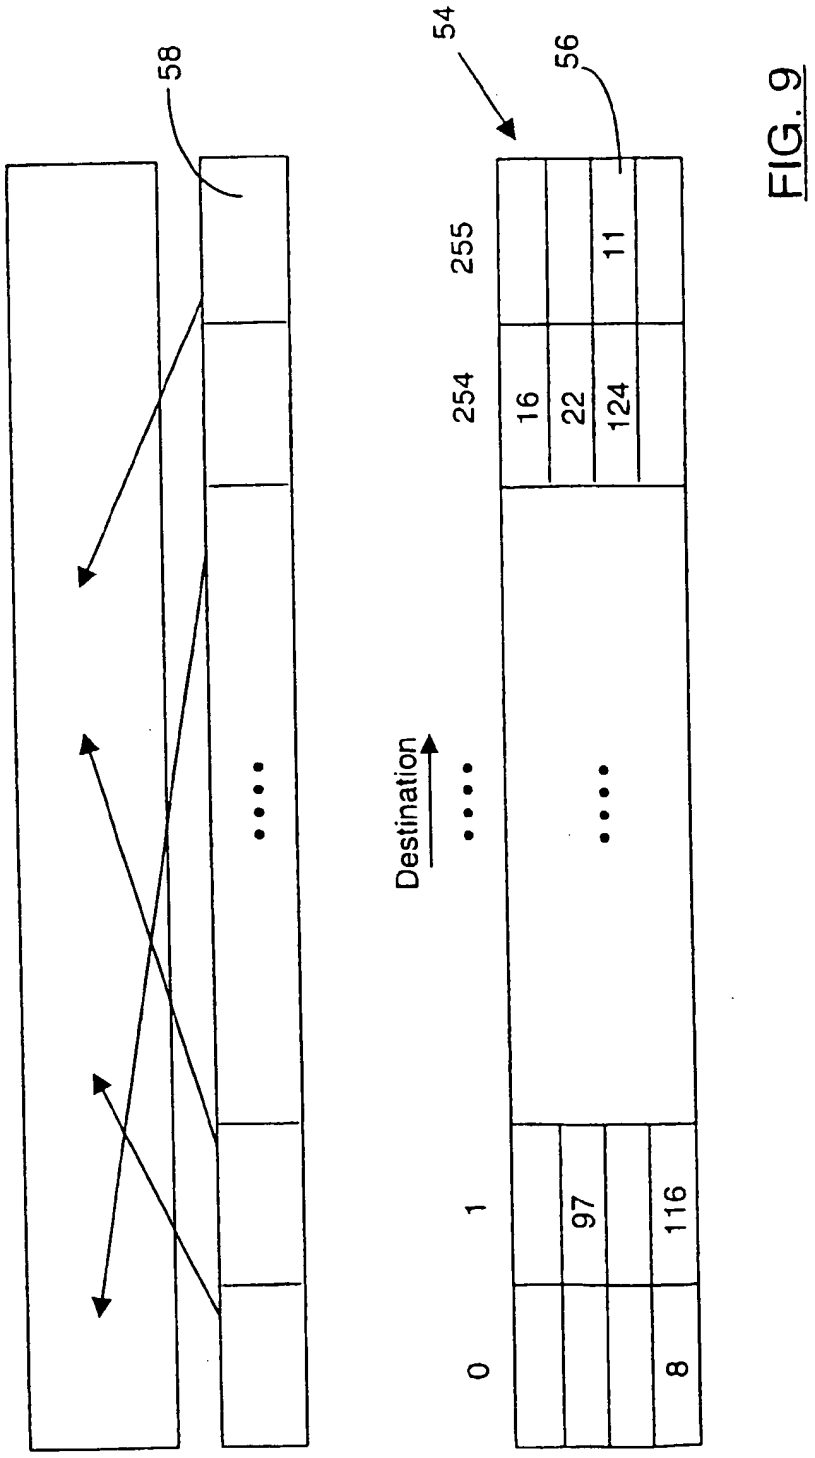

FIG. 9

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 087 635 A3

(12) EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
15.01.2003 Bulletin 2003/03

(43) Date of publication A2:
28.03.2001 Bulletin 2001/13

(21) Application number: **00308473.8**

(22) Date of filing: **27.09.2000**

(51) Int Cl.[7]: **H04Q 11/00**, H04Q 11/04, H04L 12/56

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **27.09.1999 CA 283627**

(71) Applicant: **NORTEL NETWORKS CORPORATION**
Montreal Quebec H2Y 3Y4 (CA)

(72) Inventor: **Beshai, Maged E.**
Stittsville, Ontario K2S 1E2 (CA)

(74) Representative: **Ertl, Nicholas Justin**
Elkington and Fife,
Prospect House,
8 Pembroke Road
Sevenoaks, Kent TN13 1XR (GB)

(54) High-capacity WDM-TDM packet switch

(57) A self-configuring distributed packet switch(20) that operates in wavelength division multiplexed (WDM) and time division multiplexed (TDM) modes is described. The switch comprises a distributed channel switching core (26), that includes core modules (34) which are respectively connected by a plurality of channels to a plurality of high-capacity packet switch edge modules (22,24). Each core module operates independently to schedule paths between edge modules, and reconfigures the paths in response to dynamic changes in data traffic loads reported by the edge modules. Reconfiguration timing between the packet switch modules and the channel switch core modules is performed to keep reconfiguration guard time minimized. The advantage is a high-capacity, load-adaptive, self-configuring switch that can be distributed to serve a large geographical area and can be scaled to hundreds of Tera bits per second to support applications that require very high bandwidth and a guaranteed quality of service.

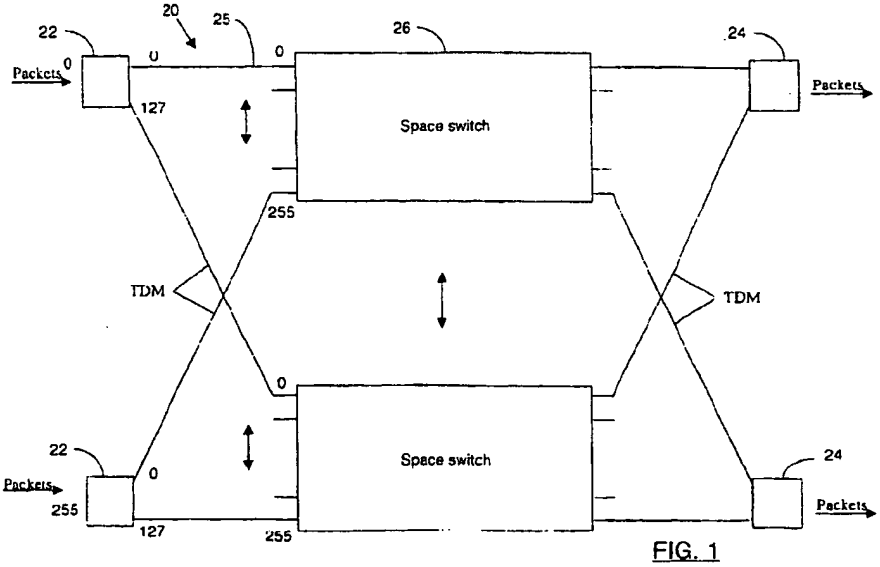

FIG. 1

EP 1 087 635 A3

European Patent Office

EUROPEAN SEARCH REPORT

Application Number
EP 00 30 8473

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X A	US 5 528 406 A (PROCTOR RICHARD J ET AL) 18 June 1996 (1996-06-18) * column 2, line 22-25 * * column 2, line 47-56 * * column 3, line 5-7 * * column 5, line 4-12 * * claim 3 * * figure 6 *	1-3,7, 10,12 13	H04Q11/00 H04Q11/04 H04L12/56
D,A	EP 0 857 383 A (NORTHERN TELECOM LTD) 12 August 1998 (1998-08-12) * the whole document *	1,13	
A	US 5 091 905 A (AMADA EIICHI) 25 February 1992 (1992-02-25) * the whole document *	1,13	

TECHNICAL FIELDS SEARCHED (Int.Cl.7)

H04Q
H04L

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
THE HAGUE	21 November 2002	Dhondt, E

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P04C01)

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO. EP 00 30 8473

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

21-11-2002

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5528406 A	18-06-1996	GB 2258360 A	03-02-1993
		EP 0551484 A1	21-07-1993
		WO 9303557 A1	18-02-1993
		JP 6501830 T	24-02-1994
EP 0857383 A	12-08-1998	US 5745486 A	28-04-1998
		DE 69619843 D1	18-04-2002
		EP 0857383 A1	12-08-1998
		JP 10512422 T	24-11-1998
		JP 3074353 B2	07-08-2000
		CA 2233628 A1	01-05-1997
		WO 9716004 A1	01-05-1997
US 5091905 A	25-02-1992	JP 1238248 A	22-09-1989
		JP 2595025 B2	26-03-1997
		CA 1317360 A1	04-05-1993
		DE 3908838 A1	05-10-1989